Co-adaptive learning over a countable space

Michael Rabadi MRABADI@SPOTIFY.COM
Spotify, Inc.
New York, NY, USA

Editor:

arXiv:1611.09816v2 [stat.ML] 30 Nov 2016

Abstract

Co-adaptation is a special form of on-line learning where an algorithm $\mathcal{A}$ must assist an unknown algorithm $\mathcal{B}$ to perform some task. This is a general framework and has applications in recommendation systems, search, education, and much more. Today, the most common use of co-adaptive algorithms is in brain-computer interfacing (BCI), where algorithms help patients gain and maintain control over prosthetic devices. While previous studies have shown strong empirical results Kowalski et al. (2013); Orsborn et al. (2014) or have been studied in specific examples Merel et al. (2013, 2015), there is no general analysis of the co-adaptive learning problem. Here we will study the co-adaptive learning problem in the online, closed-loop setting. We will prove that, with high probability, co-adaptive learning is guaranteed to outperform learning with a fixed decoder as long as a particular condition is met.

Keywords: co-adaptation, online learning, regret minimization

1. Introduction

There are many circumstances where an algorithm $\mathcal{B}$ must learn to perform a task while operating through some unknown algorithm $\mathcal{A}$. One common example is search: a person $\mathcal{B}$ must learn how to provide queries to a search engine $\mathcal{A}$ in order to get good search results. In older search engines, $\mathcal{A}$ was a fixed function. These days $\mathcal{A}$ tries to help $\mathcal{B}$ find better results using a history of $\mathcal{B}$'s queries and selections. While this isn't traditionally thought of as co-adaptive learning, experience shows that both the user $\mathcal{B}$ and the search engine $\mathcal{A}$ learn to work together to provide $\mathcal{B}$ with good results.

Co-adaptation is more explicitly studied in the context of brain-computer interfacing (BCI). BCIs are systems that help paralyzed patients gain control over a computer or a prosthetic by decoding their brain activity. In older experiments, it has been shown that patients $\mathcal{B}$ were able to learn to perform a task through a fixed decoder $\mathcal{A}$, which mapped a patient's brain signals to a computer cursor Carmena et al. (2003); Ganguly and Carmena (2009). More recently, some experimental and theoretical frameworks have been used to examine how decoders should adapt in order to improve performance DiGiovanna et al. (2009); Kowalski et al. (2013); Dangi et al. (2014). However, most of these previous studies focused on how $\mathcal{A}$ should learn without considering how $\mathcal{B}$ might adapt. A notable exception is Merel et al. (2015), where, using some strong assumptions, they were able to show that co-adaptive learning was at least as good as learning with a fixed decoder. Unfortunately, they had to assume that $\mathcal{B}$'s loss function was known to $\mathcal{A}$ and, in particular, was the mean

squared loss. Those assumptions may be reasonable for some BCI tasks, but do not hold in general.

Here we analyze the co-adaptation learning problem, which is summarized in Figure 1. In this setting, $\mathcal{B}$ has some intention y_t at time t. $\mathcal{B}$ selects an encoder $g_t \in \mathcal{G}$ with the intention of making the input $\hat{y}_{t-1}$ close to the intended output. However, $\mathcal{B}$ must work through the decoding algorithm $\mathcal{A}$, which selects a hypothesis $h_t \in \mathcal{H}$ with the intention of producing a $\hat{y}_t$ close to y_t. Previous studies have assumed that $\mathcal{A}$ has knowledge of y_t during training, or at least can infer y_t. However this is usually an unreasonable assumption. Instead, we analyze the case where $\mathcal{A}$ does not have access to y_t. A natural first question: "Is co-adaptive learning any better than learning with a fixed decoder?" We will show that, under some mild assumptions, co-adaptive learning can indeed outperform learning with a fixed decoder.

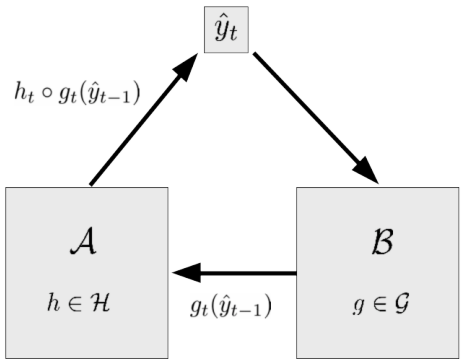

Figure 1: The co-adaptive learning scenario.

1.1 Formalization of the learning problem

Let $\mathcal{G}$ be the hypothesis set of the encoder $\mathcal{B}$, such that $\mathcal{G} = \{g : \mathcal{Y} \to \mathcal{Z}\}$, where $\mathcal{Y}$ is some countable space and $\mathcal{Z}$ is some arbitrary feature space. In practice, $\mathcal{Z}$ is typically a Hilbert space. $\mathcal{H}$ will denote the hypothesis set for the decoder $\mathcal{A}$, where $\mathcal{H} = \{h : \mathcal{Z} \to \mathcal{Y}\}$. As outlined above, we wish to find an algorithm $\mathcal{A}$ that can help another algorithm $\mathcal{B}$ minimize its loss. Specifically, $\mathcal{B}$ has a sequence of intentions $(y_1, ..., y_T)$ where $y_t \in \mathcal{Y}$. $\mathcal{B}$ encodes some control $z_t = g_t(\hat{y}_{t-1}), z_t \in \mathcal{Z}$ that it must pass through $\mathcal{A}$. $\mathcal{A}$ then selects a decoder $h_t \in \mathcal{H}$ and produces an output $\hat{y} = h_t \circ g_t(\hat{y}_{t-1})$, where $h \in \mathcal{H}$.

In the simplest case, $\mathcal{A}$ is a fixed decoder and all of the learning is done by $\mathcal{B}$, which must select a mapping $g_t \in \mathcal{G}$ at each time point t. Therefore, we wish to evaluate our co-adaptive algorithm with respect to a fixed decoder $\tilde{h} \in \mathcal{H}$. Thus, we wish to evaluate our model in terms of regret. Specifically, $\mathcal{A}$ should choose a hypothesis $h_t \in \mathcal{H}$ that minimizes its regret with respect to a fixed decoder in hindsight:

$$R_T = \sum_{t=1}^{T} L_{\mathcal{B}}(h_t \circ g_t(\hat{y}_{t-1}), y_t) - \min_{g' \in \mathcal{G}} \sum_{t=1}^{T} L_{\mathcal{B}}(\tilde{h} \circ g'_t(\tilde{y}_{t-1}), y_t)$$

where $L_{\mathcal{B}}$ is the loss function of $\mathcal{B}$ and $\tilde{y}_t$ is the predicted value of y_t when using a fixed decoder. Thus, co-adaptive learning is only better than a fixed decoder when $\sum_{t=1}^{T} L_{\mathcal{B}}(h_t \circ g_t(\hat{y}_{t-1}, y_t) < \min_{g \in \mathcal{G}} \sum_{t=1}^{T} L_{\mathcal{B}}(\tilde{h} \circ g'_t(\hat{y}_{t-1}), y_t)$ and so $R_T < 0$.

In Section 2 we will show that, under some mild assumptions, we can prove that coadaptation outperforms a fixed decoder with high probability as long as a particular condition is met.

1.2 Assumptions

Our analysis will rely on a couple of mild assumptions:

1. $\mathcal{Y}$ is a countable space.
2. $L_{\mathcal{B}}$ is Lipschitz continuous with Lipschitz constant $\ell_{\mathcal{B}} > 0$ with respect to the Hamming metric (explained below).

Assumption 1 is a straightforward assumption and applies to many problems, including search and the BCI scenario (assuming finite output). The second assumption is necessary for deriving our main theorem. We will discuss this assumption more in the next section.

1.3 Preliminaries

We will need the following theorem in order to prove the main theorem of this paper. For ease of notation, let $\psi(y) = \sum_{t=1}^{T} L_{\mathcal{B}}(\tilde{h} \circ g_t(\tilde{y}_{t-1}), y_t)$. Furthermore, set

$$\bar{\eta}_{ij} \equiv \sup_{\substack{y^{i-1} \in \mathcal{Y}^{i-1}, w, \hat{w} \in \mathcal{Y} \\ \Pr[y^i = (y^{i-1}, w)] > 0, \Pr[y^i = (y^{i-1}, \hat{w})] > 0}} \| \mathcal{L}(Y_j^n | Y^i = (y^{i-1}, w)) - \mathcal{L}(Y_j^n | Y^i = (y^{i-1}, \hat{w})) \|_{\mathrm{TV}},$$

where $\|A - B\|_{\mathrm{TV}}$ is the total variation distance between probability distributions A and B and where $\mathcal{L}(Z|Y = y)$ is the conditional distribution of Z given $Y = y$. We will then define M_T to be the largest sum of variational distances over all sub-sequences in $\mathcal{Y}^T$:

$$M_T \equiv \max_{1 \le t \le T} (1 + \bar{\eta}_{t,t+1} + \cdots + \bar{\eta}_{t,T})$$

We next define the Hamming metric $d_H(x, y) = \sum_{t=1}^{T} 1_{x_t \neq y_t}$, where $1_{x_t \neq y_t}$ is the indicator function. Thus, the Hamming metric is a count of all points that differ between two sequences. In our scenario, the Hamming metric is used between sequences y and $\tilde{y}$. We only need to evaluate the sequences with respect to $L_{\mathcal{B}}$, where $L_{\mathcal{B}}$ is $\ell_{\mathcal{B}}$-Lipschitz with respect to the Hamming metric. Specifically, we require that $|\psi(S) - \psi(S')| \le \ell_{\mathcal{B}} d_H(S, S')$, where S and S' are two sequences that differ by a single point.

With these assumptions, we can give Theorem 1 from Kontorovich et al. (2008).

Theorem 1 (Kontorovich et al. (2008) 1.1) *Without loss of generality, let $y = (y_1, ...y_T) \in \mathcal{Y}^T$ be a sequence of random variables, and $\psi : \mathcal{Y}^T \to \mathbb{R}$ is a $\ell_{\mathcal{B}}$-Lipschitz continuous function with respect to $d_H(\cdot, \cdot)$. Then, for any $\epsilon > 0$:*

$$\Pr[|\psi - \mathbb{E}[\psi]| > \epsilon] \leq 2\exp\left(\frac{-\epsilon^2}{2T\ell_{\mathcal{B}}^2 M_T^2}\right)$$

Theorem 1 is a version of McDiarmid's inequality over dependent random variables, which relaxes the i.i.d. assumption. We will use this theorem to give a probabilistic bound for the learning problem, which we will give and prove in the next section.

2. Theoretical analysis of the co-adaptive learning problem

In this section we give the main theorem, which gives a condition that guarantees that co-adaptive learning will be better than learning with a fixed decoder with high probability. Recall that we will evaluate our algorithm with respect to regret:

$$R_T = \sum_{t=1}^{T} L_{\mathcal{B}}(h_t \circ g_t(\hat{y}_{t-1}), y_t) - \min_{g' \in \mathcal{G}} \sum_{t=1}^{T} L_{\mathcal{B}}(\tilde{h} \circ g'_t(\tilde{y}_{t-1}), y_t)$$

Specifically, we can say that co-adaptive learning outperforms learning with a fixed decoder when $R_T < 0$.

Theorem 2 *Given the assumptions from Section 1.2 hold. Co-adaptation outperforms learning with a fixed decoder with probability at least $1 - \delta$ if the following holds:*

$$\sum_{t=1}^{T} L_{\mathcal{B}}(h_t \circ g_t(\hat{y}_{t-1}), y_t) + \ell_{\mathcal{B}} M_T \sqrt{2T \log \frac{2}{\delta}} < \sum_{t=1}^{T} \varepsilon_t$$

where ε_t is the expected minimal loss at time t.

Proof We want a probabilistic bound for $\psi \equiv \sum_{t=1}^{T} L_{\mathcal{B}}(\tilde{h} \circ g'_t(\tilde{y}_{t-1}), y_t)$. We can immediately derive the following bound from Theorem 1.

$$\Pr\left[\Big|\, \mathbb{E}[\psi] - \psi \Big| > \epsilon\right] \leq 2\exp\left(\frac{-\epsilon^2}{2T\ell_{\mathcal{B}}^2 M_T^2}\right)$$

Thus, with probability at least $1 - \delta$, the following holds:

$$\Big|\, \mathbb{E}[\psi] - \psi \Big| \leq \ell_{\mathcal{B}} M_T \sqrt{2T \log \frac{2}{\delta}}$$
$$\implies \psi \geq \mathbb{E}[\psi] - \ell_{\mathcal{B}} M_T \sqrt{2T \log \frac{2}{\delta}}$$

We then take the $\inf_{g' \in \mathcal{G}} \psi$, where g' is a sequence of g'_t.

$$
\begin{aligned}
\inf_{g' \in \mathcal{G}} \psi &\geq \inf_{g' \in \mathcal{G}} \mathbb{E}[\psi] - \ell_{\mathcal{B}} M_T \sqrt{2T \log \frac{2}{\delta}} \\
&= \inf_{g' \in \mathcal{G}} \mathbb{E}\left[\sum_{t=1}^{T} L_{\mathcal{B}}(\tilde{h} \circ g'_t(\tilde{y}_{t-1}), y_t) \right] - \ell_{\mathcal{B}} M_T \sqrt{2T \log \frac{2}{\delta}} \\
&= \inf_{g' \in \mathcal{G}} \sum_{t=1}^{T} \mathbb{E}[L_{\mathcal{B}}(\tilde{h} \circ g'_t(\tilde{y}_{t-1}), y_t) | Z] - \ell_{\mathcal{B}} M_T \sqrt{2T \log \frac{2}{\delta}} \\
&\geq \sum_{t=1}^{T} \inf_{g'_t \in \mathcal{G}} \mathbb{E}[L_{\mathcal{B}}(\tilde{h} \circ g'_t(\tilde{y}_{t-1}), y_t) | Z] - \ell_{\mathcal{B}} M_T \sqrt{2T \log \frac{2}{\delta}} \\
&= \sum_{t=1}^{T} \inf_{g'_t \in \mathcal{G}} \sum_{\tilde{y}_t \in \mathcal{Y}} \sum_{y \in \mathcal{Y}} L_{\mathcal{B}}(\tilde{y}_t, y_t) \mathcal{L}\big(\tilde{y}_t, y_t | Z\big) - \ell_{\mathcal{B}} M_T \sqrt{2T \log \frac{2}{\delta}} \\
&= \sum_{t=1}^{T} \varepsilon_t - \ell_{\mathcal{B}} M_T \sqrt{2T \log \frac{2}{\delta}}
\end{aligned}
$$

where $Z = \Big(\{g'_{t-i}\}_{i \in [1,t]}, \{y_{t-i}\}_{i \in [1,t]}, \{\tilde{y}_{t-i}\}_{i \in [1,t]}, \tilde{h} \Big)$, $\mathcal{L}(\tilde{y}_t, y_t | Z)$ is the joint distribution function of $(\tilde{y}_t, y_t)$ conditioned on Z, and $\{z_{t-i}\}_{i \in [1,t]}$ is the sequence of previous points for some variable z, which is all that $\mathcal{B}$ has to learn with and, consequently, use to select some $g'_t \in \mathcal{G}$.

The second equality holds by the conditional independence of the loss. The second inequality holds by the super-additivity of the min function. The third equality is the definition of the expectation of the loss function, which is a function over the joint distribution of $\hat{y}$ and y, where $\hat{y}$ is determined by selecting some $g'_t \in \mathcal{G}$. The final equality is the definition of ε_t. We can then use this lower bound on $\inf_{g' \in \mathcal{G}} \psi$ to impose the probabilistic constraint above. ■

Theorem 2 gives us a constraint that guarantees that the co-adaptive learner will outperform learning with a fixed decoder. The bound depends on properties of the conditional probabilities for the different sequences and the loss function used. The benefit of this bound is that we can use the empirical loss, $\sum_{t=1}^{T} L_{\mathcal{B}}(h_t \circ g_t(\tilde{y}_{t-1}), y_t)$ to guarantee that co-adaptive learning is better than learning with a fixed decoder with high probability. While, $\sum_{t=1}^{T} \varepsilon_t$ is not accessible in general, there may be cases where we can guarantee that, in expectation, there is some infimal error. This is certainly the case when either $\tilde{h}$ or $\mathcal{G}$ is not rich enough to perform some task with zero error. Indeed, it is easy to see that for all $\tilde{h}$ and $\mathcal{G}$, there exists some set of conditional distributions $\mathcal{L}(\tilde{y}_t, y_t | Z), \forall \tilde{y}, y \in \mathcal{Y}^T$ where the infimal error is greater than zero. Indeed this occurs whenever $\mathcal{L}(\tilde{y}_t = y_t | y_t, Z) < 1$ or $\mathcal{L}(y_t | Z) < 1$, which holds by the definition of the conditional joint distribution:

$$\mathcal{L}(\tilde{y}_t, y_t | Z) = \mathcal{L}(\tilde{y}_t | y_t, Z) \mathcal{L}(y_t | Z)$$

3. Conclusion

We gave a condition for the co-adaptive learning problem that, if verified, guarantees that co-adaptive learning outperforms learning with a fixed decoder with high probability. This work marks a significant generalization of the co-adaptive learning problem. Our analysis is more general than one with the standard i.i.d. assumption. However, one drawback is that the derived constraint is not accessible in general: we do not always know the expected minimal loss as a function of time for $\mathcal{B}$ learning over a fixed decoder. However, in practice we may be able to approximate ε_t by measuring the learning rate and initial transition probability.

References

Jose M Carmena, Mikhail A Lebedev, Roy E Crist, Joseph E O'Doherty, David M Santucci, Dragan F Dimitrov, Parag G Patil, Craig S Henriquez, and Miguel AL Nicolelis. Learning to control a brain–machine interface for reaching and grasping by primates. *PLoS biol*, 1 (2):e42, 2003.

Siddharth Dangi, Suraj Gowda, Helene G Moorman, Amy L Orsborn, Kelvin So, Maryam Shanechi, and Jose M Carmena. Continuous closed-loop decoder adaptation with a recursive maximum likelihood algorithm allows for rapid performance acquisition in brain-machine interfaces. *Neural computation*, 2014.

Jack DiGiovanna, Babak Mahmoudi, Jose Fortes, Jose C Principe, and Justin C Sanchez. Coadaptive brain–machine interface via reinforcement learning. *IEEE transactions on biomedical engineering*, 56(1):54–64, 2009.

Karunesh Ganguly and Jose M Carmena. Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol*, 7(7):e1000153, 2009.

Leonid Aryeh Kontorovich, Kavita Ramanan, et al. Concentration inequalities for dependent random variables via the martingale method. *The Annals of Probability*, 36(6): 2126–2158, 2008.

Kevin C Kowalski, Bryan D He, and Lakshminarayan Srinivasan. Dynamic analysis of naive adaptive brain-machine interfaces. *Neural Computation*, 25(9):2373–2420, 2013.

Josh Merel, Donald M Pianto, John P Cunningham, and Liam Paninski. Encoder-decoder optimization for brain-computer interfaces. *PLoS Comput Biol*, 11(6):e1004288, 2015.

Josh S Merel, Roy Fox, Tony Jebara, and Liam Paninski. A multi-agent control framework for co-adaptation in brain-computer interfaces. In *Advances in Neural Information Processing Systems*, pages 2841–2849, 2013.

Amy L Orsborn, Helene G Moorman, Simon A Overduin, Maryam M Shanechi, Dragan F Dimitrov, and Jose M Carmena. Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron*, 82(6):1380–1393, 2014.